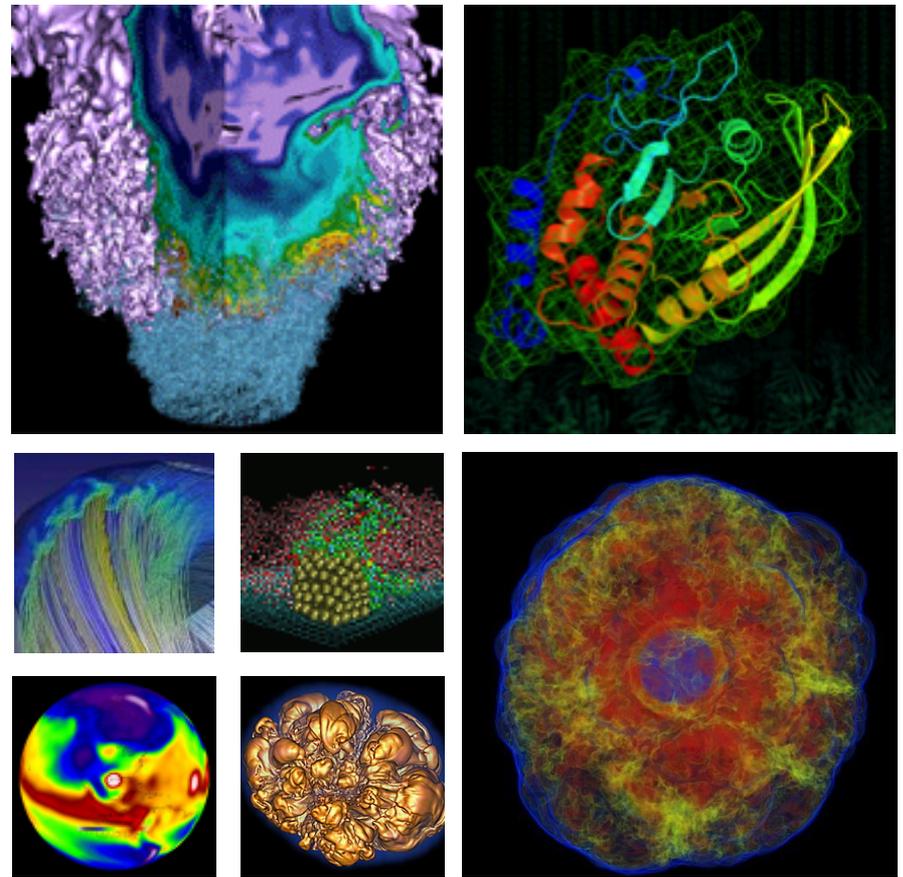


Towards a Light Source Data and Analysis Framework at NERSC



Jack Deslippe, Shane Canon,
Eli Dart, Abdelilah Essiari,
Alexander Hexemer, Dula
Parkinson, Simon Patton,
Craig Tull + Many More



The ALS Data Needs

Light source data volumes are growing many times faster than Moore's law.

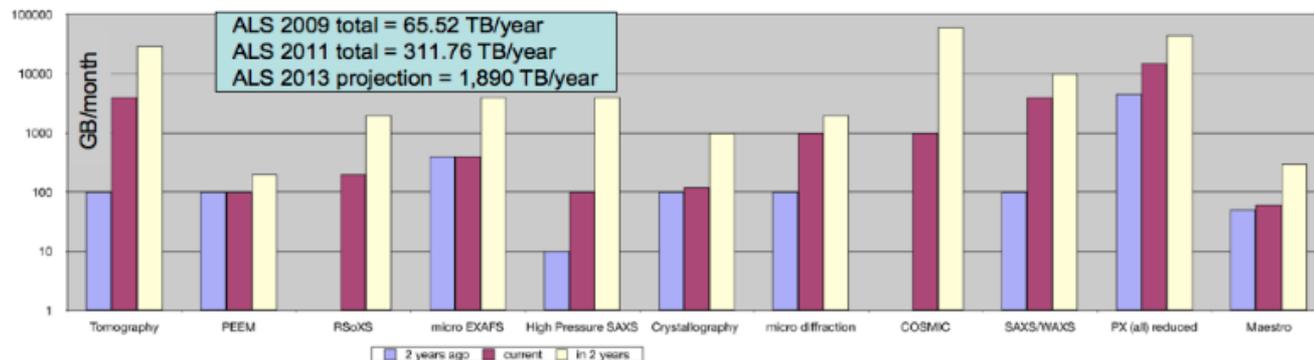
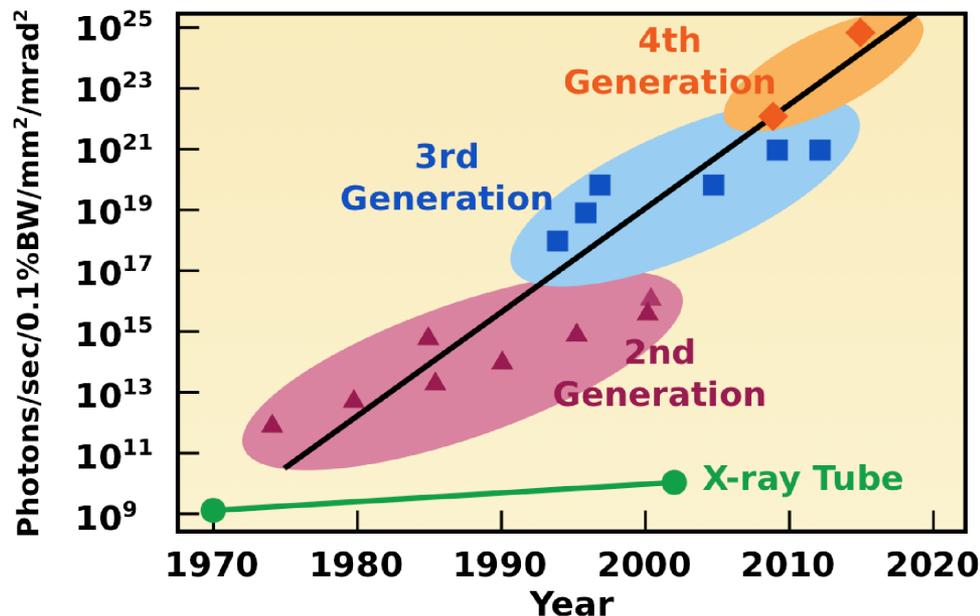
- Light source luminosity
- Detector resolution & rep-rates
- Sample automation

BES user facilities serve 10,000 scientists and engineers every year.

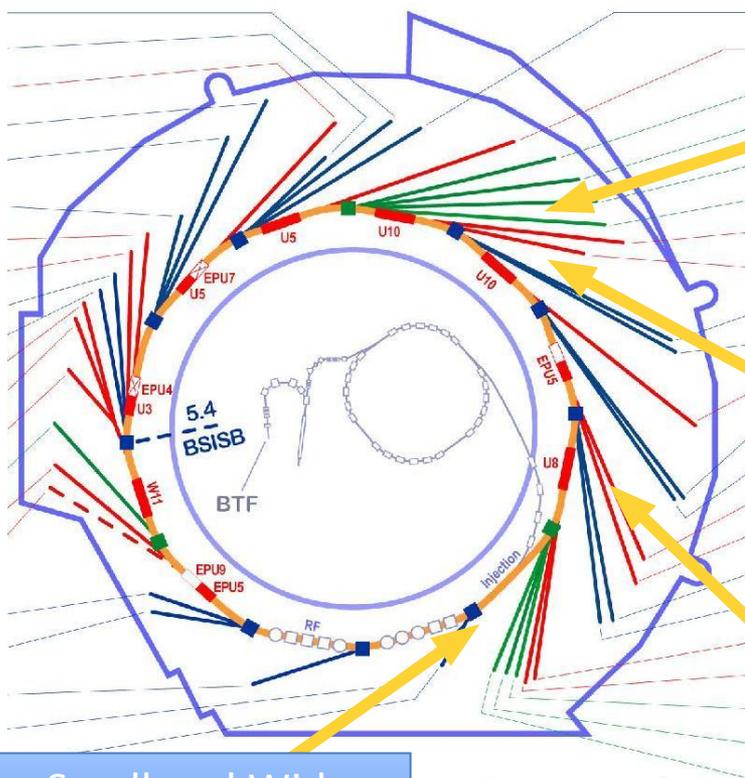
Mostly composed of many small groups with no SW experts.

Simulation & analysis traditionally done on very small scale resources.

BES software landscape is largely ad-hoc and relies heavily on a limited number of experts to handle analysis.

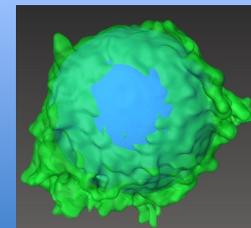


Diversity at the ALS



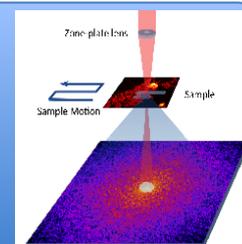
Tomography
(Real Space)

Arec3d, QuantCT,
CrunchFlow



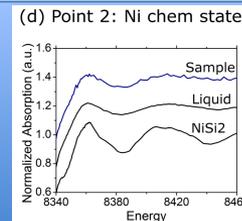
COSMIC Hybrid
(Real and
Reciprocal Space)

Ptychographic
reconstruction



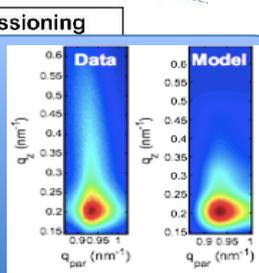
MicroXAS
(Spectroscopy)

ShirleyXAS (MSD)
BerkeleyGW (NERSC)

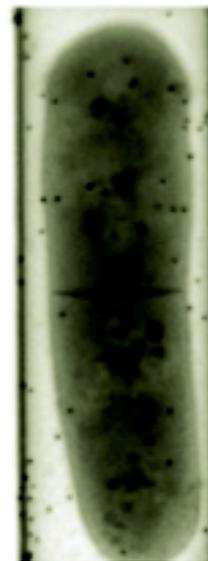
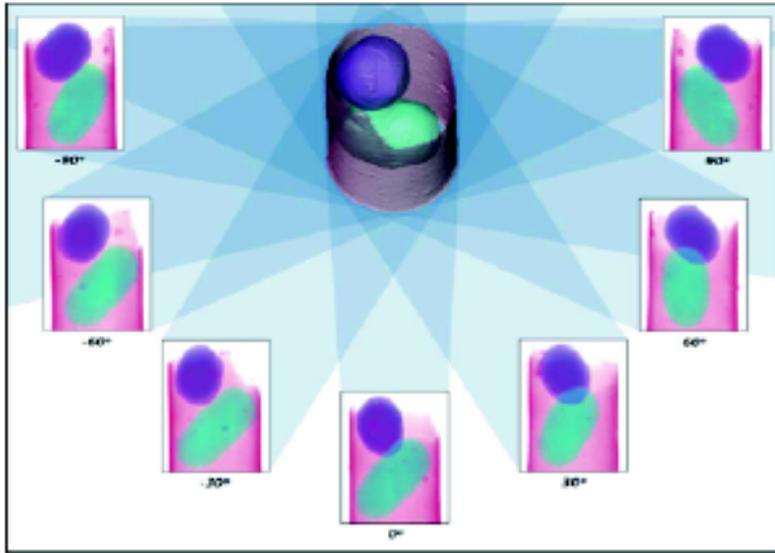


Small and Wide
Angle Scattering
(Reciprocal Space)

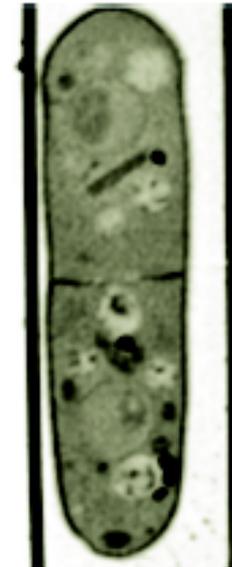
HipGISAXS/HipMC
parallel Scattering



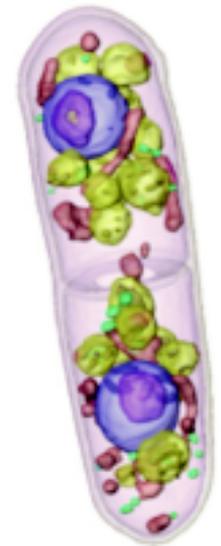
Tomography Beamline 8.3.2



Raw



Normalized



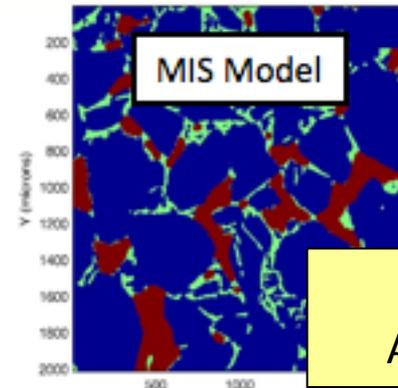
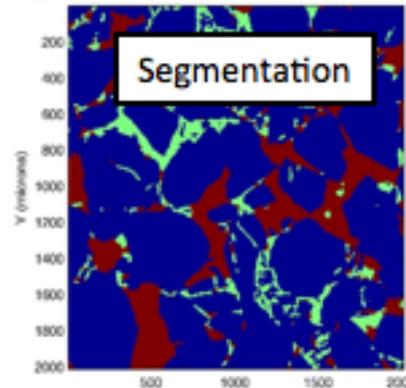
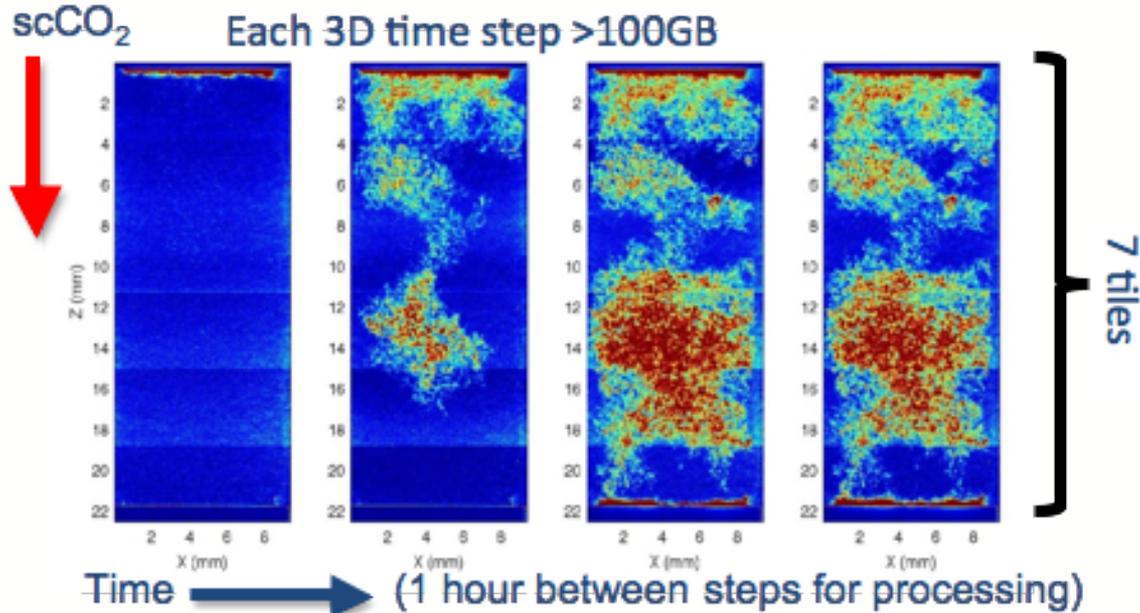
Reconstructed

Measure X-Ray loss through a sample at various angles.
"Reconstruct" internal structure of sample from images.

Traditionally most users don't achieve a reconstruction their data.

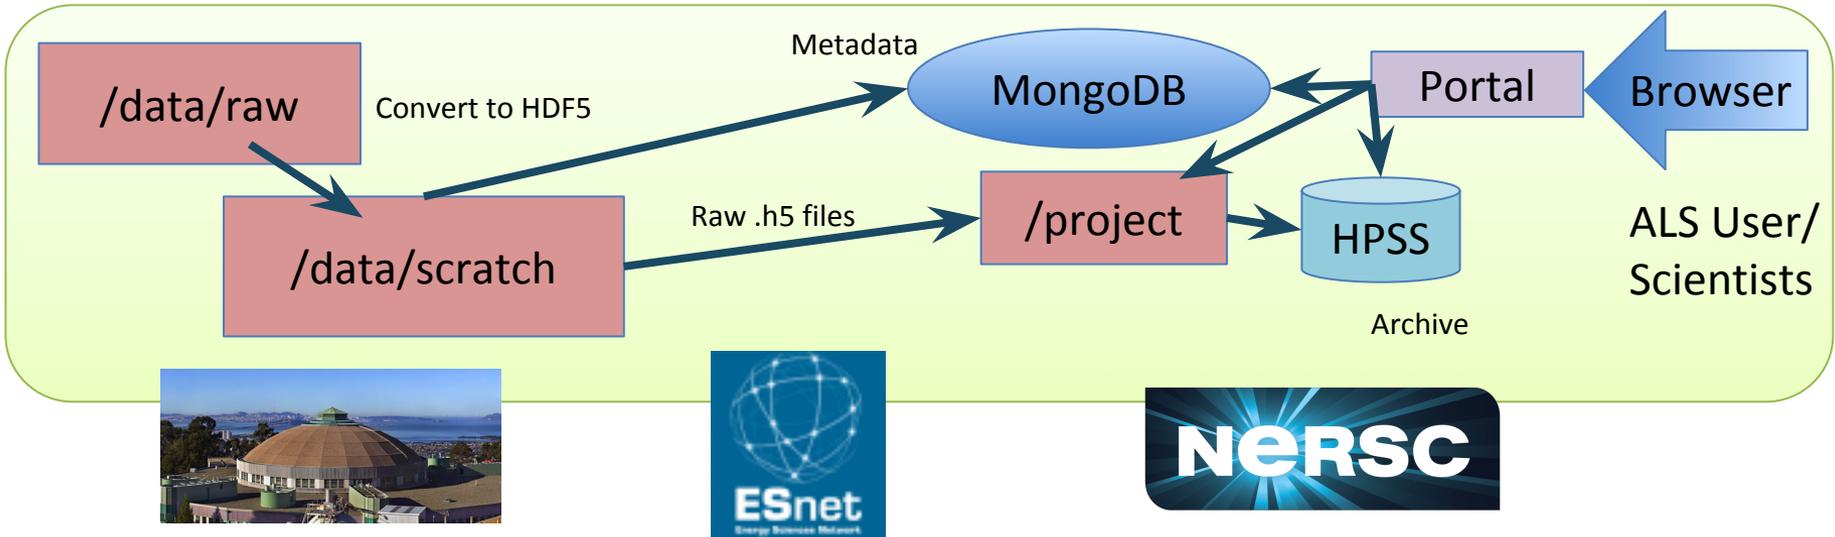
Geologic CO₂ Sequestration

- 3-D beam-time feedback currently not available. Many users need this capability.
- Offline processing and simulation of large data sets impossible for many users.
- 3-D reconstruction often requires expert's direct intervention.

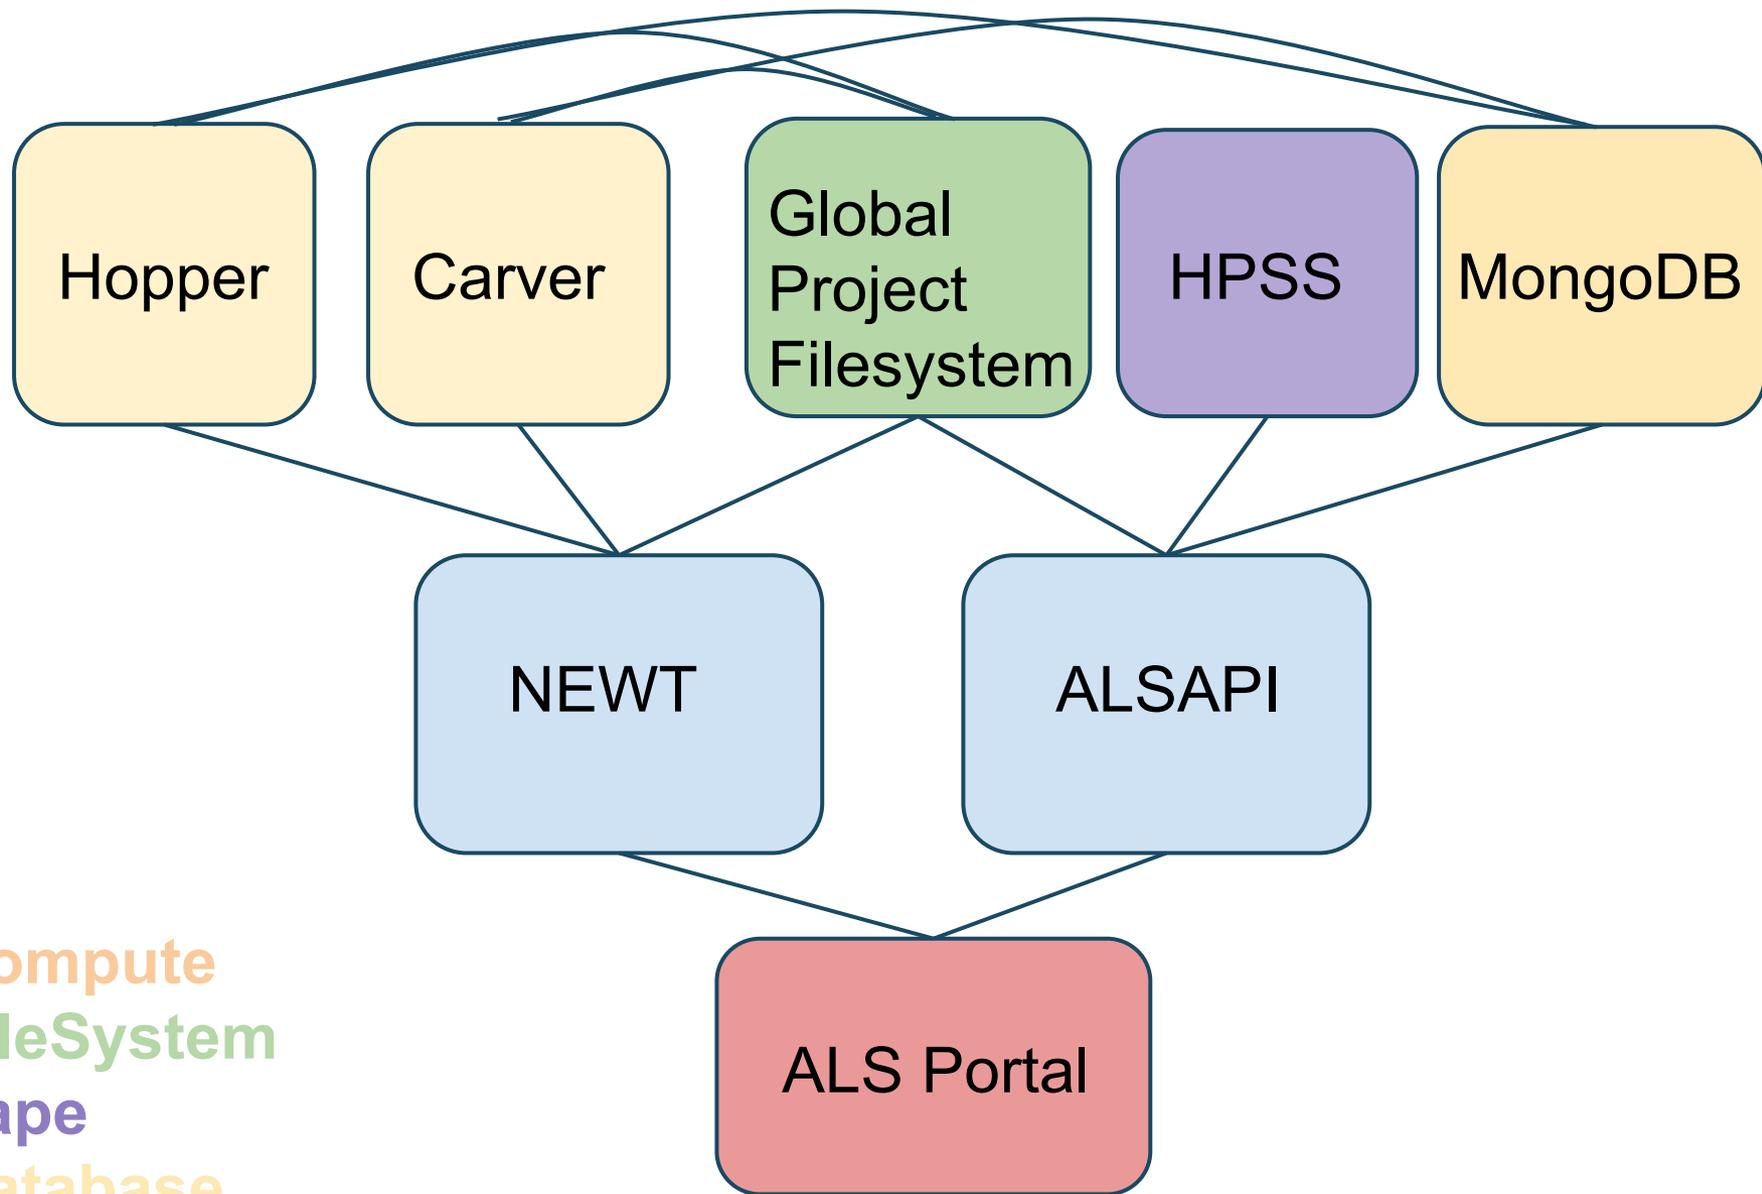


Jonathan
Ajo-Franklin

Data Pipeline Prototype in Action



- Prototype ALS Pipeline
 - Package Data into HDF format, existing data model standard (Nexus)
 - Capture critical metadata in HDF and Schema-less DB (MongoDB)
 - Automatically transfer to NERSC
 - Prototype web-access to allow ALS users to access data
- Investigating Methods to Store and Query Raw Sequence Data
 - Numbers of elements and query patterns exceed capabilities of existing off the shelf technologies.



Hopper

Carver

Global
Project
Filesystem

HPSS

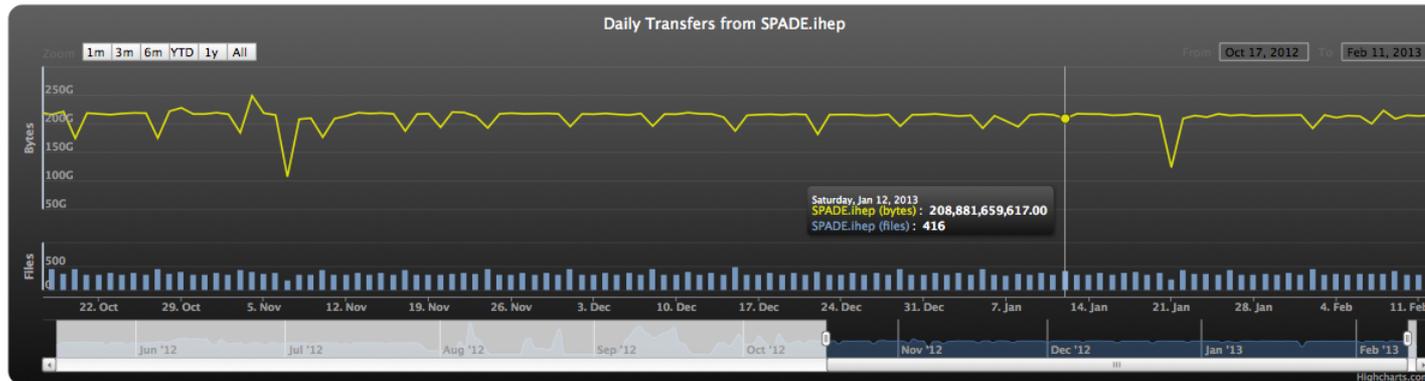
MongoDB

NEWT

ALSAPI

ALS Portal

Data Transfer Pipeline



<http://nest.lbl.gov/projects/spade/html/>

- ALS Data Transfer Node (DTN): ESNNet reference architecture
- Spade – Dayabay/IceCube production tool.
- Reliably transfer data from source to warehouse.
- Trigger actions at warehouse.
- Validate cksum and release source data file.
- GridFTP, Globus-Online, or scp transfer protocols
- Web and Mobile based monitoring and control

HDF5 Advantages:

- Can compactly store multiple images in an organized way
- Can attach metadata to various datasets
- Highly performant IO for analysis

MongoDB Advantages:

- Flexible document store; different document fields for different beamlines
- Can create rich document-subdocument structure
- Can add/remove fields without worrying about schema
- Can index across multiple fields
- Outputs to JSON

The APIs



The "ALSAPI"

DJANGO App, runs on portal-auth.nersc.gov as "alsdata" user. Stages, converts, archives, retrieves and catalogs ALS data sets. Same "auth" as NEWT.

NEWT

Users authenticate with NERSC credentials. Can run analysis and simulation software on Hopper and Carver. Results saved in user's area with option to share.



The Interface Features



Welcome to the ALS Data and Simulation Portal

This site is a very large work in progress. Log in with your NERSC account at the top right to gain access to NERSC compute resources (required for simulation and analysis tools).

[Learn more. >](#)



- What is this?
- [Search Database](#)
- [Running Calculations](#)
- [Finished Calculations](#)
- [Submit Calculations](#)
- [Draw Molecule](#)

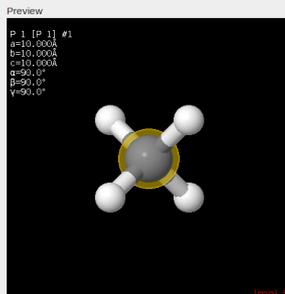
Submit Calculation

Material Name

Excited Atoms

Coordinates

C -0.00025 -0.00025 -0.00025
 H 0.64018 0.64018 0.64018
 H -0.64075 -0.64075 0.64049
 H -0.64075 0.64049 -0.64075
 H 0.64049 -0.64075 -0.64075

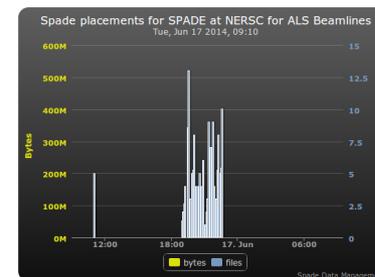


Upload Coordinates
 No file chosen

Current Disk Usage



Data Transfer Status



Data Browser

Enter Search Term

Sort By:

e438-r0138-s04-c00.h5

e438-r0140-s02-c00.h5

e438-r0139-s04-c00.h5

e438-r0139-s02-c00.h5

e438-r0137-s04-c00.h5

e438-r0136-s04-c00.h5

e438-r0138-s02-c00.h5

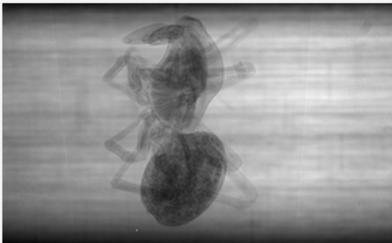
e438-r0140-s00-c00.h5

The Interface Features

Data Browser Select All Deselect All

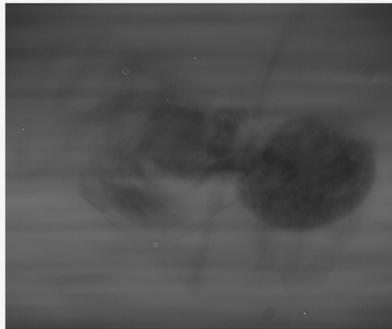
chilarch ? Sort By: Processed Date Descending

20130713_185717_Chilarchaea_quellon_F_9053427_IKI_h5



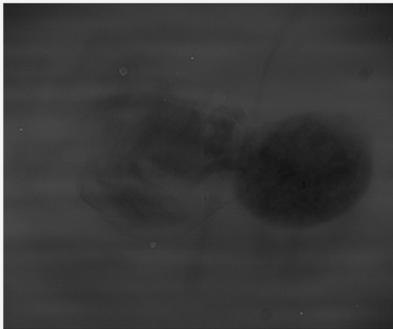
on tape
 imgrec gridrec norm available

20130524_173053_Chilarch_9053427_juv_IKI_33-5keV_h5



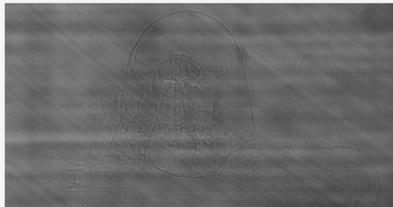
on disk
 norm norm imgrec gridrec available

20130524_164630_Chilarch_9053427_juv_IKI_h5



on disk
 norm imgrec gridrec norm available

20130426_134239_SYN_2_9_Chilarchaea_F_9053427_HW0078.h5



on disk
 norm imgrec gridrec gridrec norm available

The Interface Features

ALS About Status & Tools Data Browser Simulation Welcome jdeslip Logout

raw norm sino gridrec imgrec raw images Download H5 File



Basic Utilities

img0751_thumb.jpg

False Color: Gray Jet Sine

Smoothing

Histogram

Level: 0

Window: 255

Navigation Keys

facility: als
end_station: bl832
sdate: 2013-07-13 18:57:27
owner: hmwood
experimenter:
senergy: 33501.088010
pxsize: 0.001300
nangles: 2049
nslices: 1596
nrays: 2560
arange: 180.000000
usebrightexpose: 0
cooler_target: 0.000000
stepdeg: 0.000000
xtechdir: c:\ct\standards
cdmaterial: LEAD
noisesig: 0.001500
rsize: 0.000000
cammode: 0

The Interface Features

The screenshot displays the NERSC interface with a central 'Volume Render' window. The window title is 'Volume Render' and it contains a 3D visualization of two biological structures, possibly cells or molecules, rendered in a purple and blue color scheme. Below the visualization are control buttons: '-', '+', 's', 'S', 'HiQ', 'on', and a dropdown menu. To the right of these buttons, it says 'Steps: 192'. Below the buttons are three horizontal sliders for color mapping. The top slider is red and ranges from 0 to 255. The middle slider is green and ranges from 0 to 255. The bottom slider is blue and ranges from 0 to 255. To the right of the Volume Render window is a metadata panel with the following information:

- facility: als
- end_station: bl832
- sdate: 2013-07-13 18:57:27
- owner: hmwood
- experimenter:
- senergy: 33501.088010
- pxsize: 0.001300
- nangles: 2049
- nslices: 1596
- nrays: 2560
- arange: 180.000000
- usebrihtexpose: 0
- cooler_target: 0.000000
- stepdeg: 0.000000
- xtechdir: c:\ct\standards
- cdmaterial: LEAD
- noisesig: 0.001500
- rszsize: 0.000000
- cammode: 0

Challenges and Future Directions



- Supporting multiple beamlines is a scaling challenge. Abstracting workflow, data transfer and providing a composable analysis framework is necessary
- How guarantee real-time analysis at center with 99% usage and multi-day backlog. Killable-Queue
- Turn-around on the order of minutes requires hiding communication costs - overlapping communication and computation.
- Standardization of data format across centers and beamlines. 3D vs 2D datasets within HDF5.
- Visualizing 100 GB datasets at full scale without jank.
- WAN distribution of data back

to end-users

Total Datasets	89024
Total Data	477.9 TB
Total Raw Data	82.0 TB
Total Data Archived on Tape	477.9 TB
Total Data Staged on Disk	152.3 TB
Users with Data in System	114
Size of Mongo Database	39.25 GB

Contributors



- Craig E. Tull (CRD), A.Hexemer (ALS), B.Tierney (SND/ESNet), D.Prendergast (MSD), J.Deslippe (NERSC)

- Lab-wide LDRD (5 Divisions) funded FY 2013

Computing Research Division, Advanced Light Source, Scientific Networking Division, Molecular Foundry, NERSC

- Richard Shane Canon [NERSC], Elaine Chan [ALS], Eli Dart [ESNet], Abedelilah Essiari [CRD], Dan Gunter [NERSC], Alexander Hexemer [ALS], Xiaoye Sherry Li [CRD], Filipe Maia [NERSC], Stefano Marchesini [ALS], Matthew Marcus [ALS], Inder Monga [ESNet], Dilworth Parkinson [ALS], Simon Patton [CRD], Eric Pouyoul [ESNet], David Prendergast [MSD], Lavanya Ramakrishnan [CRD], David Skinner [NERSC]

EXTRA SLIDES